

Combined Surface and Volume Processing for Fused Joint Segmentation*

Peter R. Krekel^{*,a}, Edward R. Valstar^a, Frits H. Post^b,
P. M. Rozing^a, Charl P. Botha^b

^aDepartment of Orthopaedics, Leiden University Medical Center, Leiden, The Netherlands

^bVisualisation Department, Delft University of Technology, Delft, The Netherlands

Abstract

Purpose Segmentation of rheumatoid joints from CT images is a complicated task. The pathological state of the joint results in a non-uniform density of the bone tissue, with holes and irregularities complicating the segmentation process. For the specific case of the shoulder joint, existing segmentation techniques often fail and lead to poor results. This paper describes a novel method for the segmentation of these joints.

Methods Given a rough surface model of the shoulder, a loop that encircles the joint is extracted by calculating the minimum curvature of the surface model. The intersection points of this loop with the separate CT-slices are connected by means of a path search algorithm. Inaccurate sections are corrected by iteratively applying a Hough transform to the segmentation result.

Results As a qualitative measure we calculated the Dice coefficient and Hausdorff distances of the automatic segmentations and expert manual segmentations of CT-scans of ten severely deteriorated shoulder joints. For the humerus and scapula the median Dice coefficient was 98.9% with an interquartile range (IQR) of 95.8% - 99.4% and 98.5% (IQR 98.3% - 99.2%) respectively. The median Hausdorff distances were 3.06 mm (IQR 2.30 mm - 4.14 mm) and 3.92 mm (IQR 1.96 mm - 5.92 mm) respectively.

*Corresponding author. Tel.: +31 71 526 1566; fax: +31 71 526 6743. E-mail address: P.R.Krekel@LUMC.nl (Peter R. Krekel)

*This is a preprint of the article with the same name, published in the Springer International Journal of Computer Assisted Radiology and Surgery and appearing as epub on December 5, 2009 with DOI 10.1007/s11548-009-0400-4. The Springer version can be accessed at <http://www.springerlink.com/content/582m357922024815/>.

Conclusion The routine satisfies the criterion of our particular application to accurately segment the shoulder joint in under two minutes. We conclude that combining surface curvature, limited user interaction and iterative refinement via a Hough transform forms a satisfactory approach for the segmentation of severely damaged arthritic shoulder joints.

Key words: Segmentation, Curvature, Hough Transform, Rheumatoid Joints, Glenohumeral Joint

1. Introduction

In orthopaedic surgery, CT-scans are routinely used to diagnose and plan joint replacement surgery for severe cases of osteoarthritis and rheumatoid arthritis. These diseases are the two most common forms of arthritis and are characterised by heavy deterioration of the cartilage and bone surface of joints. The CT-scans yield anatomical information that is hard to obtain through other means. A wide variety of orthopaedic applications require CT-scans to provide pre- and intra-operative assistance to the surgeon [12, 7, 21].

Segmentation of the data enables the surgeon to distinguish the geometry of the different bones of the joint. This allows for fast assessment of the state of the pathological joint and provides means to set up a pre-operative plan for surgery. One of the target applications for our segmentation routine utilizes the bone geometry for biomechanical modelling of the shoulder. The challenging aspect of this segmentation problem is that bone tissues need to be separated, rather than just classified.

For arthritic joints the segmentation process is typically complicated by large variations in bone density and irregularities of the bone surface (see Figure 1 for an example slice of such a dataset). The varying bone density complicates interpretation of the data. In addition, the joint often appears to be fused in the CT data due to the significantly decreased joint space. As the cartilage of pathological joints wears off, the joint space can become virtually untraceable. This is further complicated by the limited resolution of the CT scanning equipment. Manual slice-by-slice segmentation is labor intensive and may take up to two hours per dataset.

Existing segmentation techniques, including a high-level integration of level-sets and watershed segmentation, have limited success with this specific segmentation task, as discussed by Botha [2]. Although these techniques lead to accurate segmentation results, they require adaptation of numerous parameters. This is a

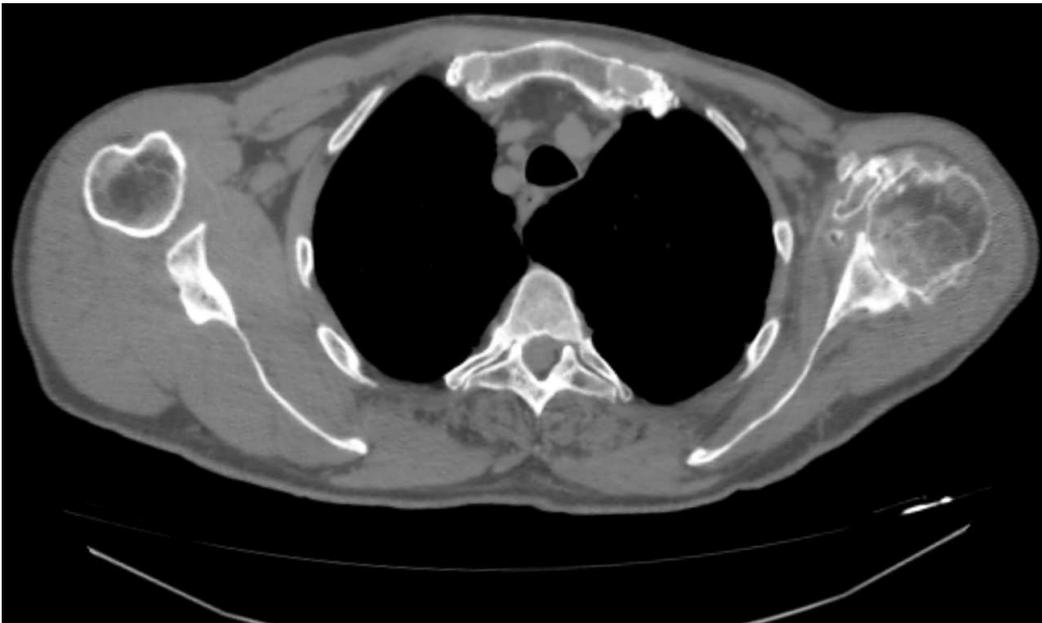


Figure 1: A single slice of a CT-dataset used in this study. The image shows the right and left shoulder in the transverse plane (i.e. view from feet to head). As can be seen, the left shoulder joint (on the right side of the image) is heavily deteriorated due to rheumatoid arthritis, complicating segmentation of this shoulder.

time-consuming process and therefore these techniques are not suitable for clinical protocols.

One of the promising techniques for separating skeletal structures is presented by Kang et al. [11]. The technique consists of 3D region growing with locally adaptive thresholds, followed by a mixture of 3D and 2D morphological operations to close holes in the segmentation. The resulting segmentation is then smoothed by adjusting its containing iso-surface. The technique is applied to the hip, the knee and the skull. The authors state that a site-specific approach is required for separating different bony structures at joints *before* their techniques can be applied and that this separation is site specific. In our opinion, the determination of this site-specific separation method is one of the most challenging tasks in the case of the shoulder due to its deformed shape.

In the work of Zoroofi et al. [29] rheumatoid hip joints are segmented by making use of histogram-based thresholding. A number of landmarks are chosen relative to the convex hull. These points are used to construct an initial ellipsoid around the femoral head. This ellipsoid is used to derive an initial estimate of the joint space, which is subsequently refined on a Hessian-filtered version of the data. Although their method yields accurate segmentation results for less arthritic joints, the joint spaces of severely arthritic and extremely arthritic hips were only correctly segmented 30% and 12% respectively.

Branzan-Albu et al. [4] and Tremblay et al. [24] describe a technique for segmenting shoulders from T1-weighted MRI images of healthy shoulders. Although this work is based on MRI data rather than CT data, it is relevant because it specifically targets the shoulder. A separate 2D segmentation is performed on each 2D slice. This segmentation is based on Wiener filtering, Sobel edge detection and region growing. Morphological techniques are used to successfully reconstruct a smooth 3D segmentation from the 2D segmentations. However, this work focuses on healthy subjects with large acromiohumeral and glenohumeral distances. Their technique does not address the problems of decreased joint space and bone decalcification. It does however indicate the promise of MRI-based techniques for the segmentation of the shoulder skeleton.

Model-based segmentation techniques can also lead to a successful solution for the segmentation of anatomical shapes [22, 18, 28]. In orthopaedics, statistical shape models for hips and knees have been investigated [20, 10]. However, the great variability in healthy shoulder bone anatomy [19, 27], together with the numerous small and unpredictable variations due to the rheumatoid state of the shoulders that we wish to segment, greatly complicate the successful application of shape models.

We were inspired by curvature-based segmentation techniques that operate directly on 3-D surface meshes. For example, by adapting the morphological watershed to segment regions of similar surface curvature on 3-D meshes, object surfaces can be decomposed into meaningful regions [15, 17]. This generally works well for the segmentation of solid objects that give a sharp contour in the image data. However, for organic structures watershed techniques lead to less satisfactory results.

To the best of our knowledge no algorithm exists that can be used for the segmentation of severely deteriorated shoulder joints from CT data. The discussed techniques require extensive adaptation to be applicable to this difficult segmentation problem. Usage of these adapted techniques would require image processing expertise and therefore they cannot be included in day-to-day clinical workflow.

In this paper we describe a new technique that exploits the spherical shape of the humeral head to automatically segment the rheumatoid joints. The center of rotation of the glenohumeral joint can be determined by applying a Hough transform directly on the CT data, as described in [25]. In the pathological case, this spherical shape may have partially disappeared and often appears to be fused with the shoulder blade. However, the remaining curved shapes give an indication of where the joint gap is located. Our technique utilises this knowledge and, with a minimum of user interaction, segments highly pathological glenohumeral joints in less than two minutes.

The major contribution of our work is that we combine curvature information of surface data with path searching techniques on volumetric data. By combining these techniques, we retain fine details of the volumetric data, enabling this fast routine to result in accurate bone models. Concepts traditionally from visualisation are combined with concepts traditionally from image processing in order to solve a challenging segmentation problem. Expert orthopaedic surgeons stated that the segmentation process is sufficiently accurate and fast to be included in a clinical protocol, enabling a wide variety of diagnosis and surgical planning purposes.

The remainder of this paper is structured as follows: Section 2 presents the details of our new segmentation technique and section 3 documents the results of an evaluation we performed on ten pathological shoulder CT datasets. In Section 4 we present our conclusions and plans for future work.

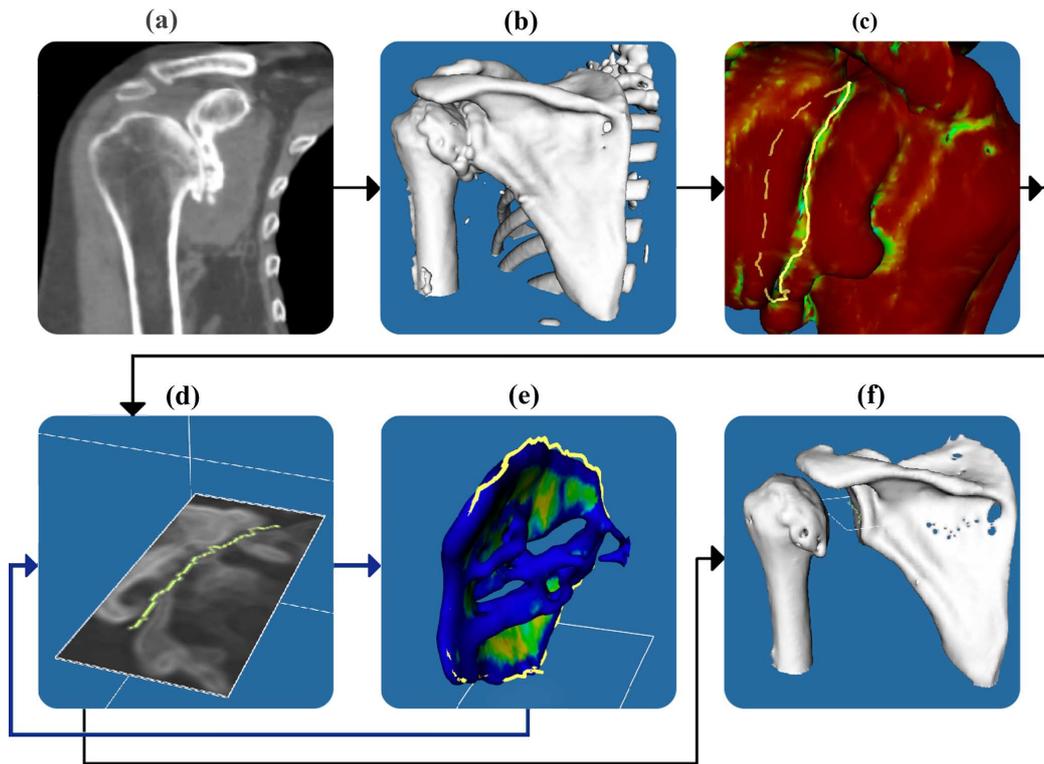


Figure 2: Overview of the stages of our segmentation method. (a) The input data. (b) The data after preprocessing, described in Section 2.1. (c) The joint separation loop, described in Section 2.2. (d) Slice-by-slice separation contours (Section 2.3). (e) The Hough feature volume (Section 2.4). (f) The final segmentation result, obtained after iteratively applying the slice-by-slice separation contours and the Hough feature volume (Section 2.5).

2. Method Description

The severity of rheumatoid arthritis can be represented by the widely accepted Larsen-score [13]. It is determined using radiographs and indicates how much bone damage has been done by the disease. A high Larsen score generally means that a corresponding CT-scan will be difficult to segment. We have developed our techniques for segmentation of the highest severity level (level 5) of this score. This level of the Larsen-score implies that the joint has deformed from its normal anatomical shape and the density of the bones is not uniformly distributed. In addition, holes are scattered over the articular surface of the bone. For many of the slices the joint space is not visible. These factors lead to a challenging segmentation problem.

Our segmentation technique consists of four stages (see Figure 2) of which the first stage is a preprocessing stage to counteract variability in the scan parameters and to alleviate noise. During the second stage a contour is determined that encircles the glenohumeral joint. The third stage connects the pairwise intersection points of this loop and the individual CT-slices. This temporary segmentation is refined during the fourth stage by creating a surface model of the connecting lines and applying a Hough transform to this model to determine how plausible the different parts of the segmentation are. The Hough transform allows us to improve the segmentation results at noisy parts of the CT-data, difficult to interpret due to the rheumatoid state of the joints.

2.1. Preprocessing

Clinical CT data is most often anisotropically sampled with the slice thickness being greater than the pixel size. To facilitate later processing steps, the data is isotropically resampled with trilinear interpolation, after which it is smoothed with a Gaussian filter with a standard deviation of 1.5 and radius of 2 voxels to counter noise. The window size is chosen so that the joint space is preserved after smoothing. The smoothed volume is used for the second stage of our approach that requires a volume with larger scale features. The unsmoothed isotropic volume that is used for the actual slice-by-slice segmentation should contain all details of the original volume. To enhance the edges of the volume, it is sharpened with the following voxel-wise operations: 1) The volume and its gradient magnitude are summed; 2) The gradient magnitude of this sum is subtracted from the volume.

The Hounsfield Scale is defined such that distilled water results in a Hounsfield Unit (HU) of 0, whereas air at Standard Temperature and Pressure results in a HU

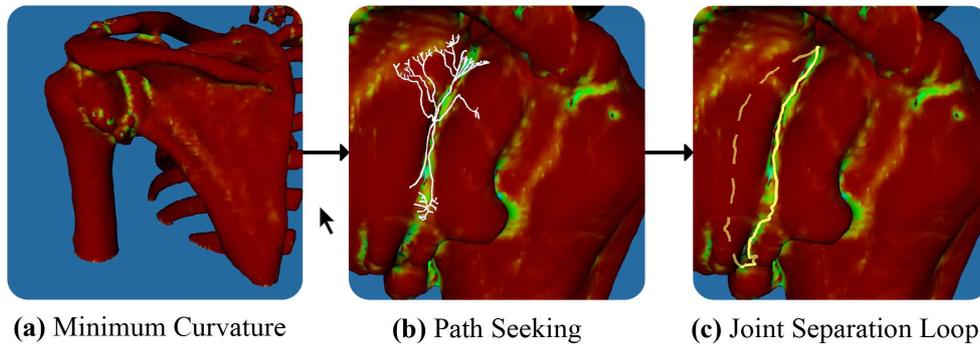


Figure 3: Extraction of the joint separation loop. (a) First, the minimum curvature of the rough surface model is calculated. Areas with negative minimum curvature are colored green. (b) The user selects a green part of the surface model where the glenohumeral joint is located and this initiates a path searching algorithm. (c) The path seeking algorithm uses a cost function to encircle the glenohumeral joint and return this edge loop as output.

of -1000. With these standards cortical bone has a HU of 400 and rheumatoid bone has a HU of 200. Surrounding muscle tissues have a HU of 50. As a result of this, the gradient magnitude of the pathological articular surface varies between 0 and 160.

2.2. Joint Separation Loop

We define the joint separation loop as the contour that encircles the fused humerus and scapula, running along the valley, an elongated surface depression, where the bones meet. The joint separation loop will be used in later stages to derive a surface that separates the humerus and scapula. In this section we document the determination of the joint separation loop. Figure 3 gives an overview of this stage.

To initiate this stage, the smoothed volume is converted to a coarse surface model of the bones using the Marching Cubes algorithm [14]. Although healthy bone has a HU of approximately 400, experiments showed that a threshold value of 100 captures most of the bone geometry when it is in a pathological state.

Due to joint space narrowing the result is a single surface model that represents a fused scapula (i.e. shoulder blade) and humerus (i.e. upper arm). The different shapes of these bones result in a clearly visible valley between the scapula and the humerus that can be extracted by finding an edge loop that runs along vertices where the minimum curvature has a great magnitude.

In order to determine the required surface curvature, based on our triangular

mesh, we have used the technique described by Page et al. [16]. Their method determines the surface curvature at each vertex of a mesh as follows: First all triangles in a neighbourhood of configurable size around the vertex are extracted. Then the normals of these triangles are combined with an inverse distance weighted averaging to determine the normal at the vertex. Finally this central normal and the positions of the vertices of the surrounding triangles are used to derive a curvature tensor, of which the Eigen-analysis yields the principal curvatures at the central vertex.

To extract the joint separation loop, the user selects any part of the glenohumeral joint. The vertex closest to the point selected by the user acts as the seed for a surface constrained region growing method. The seed is the initial region. Neighbouring vertices, i.e. nodes, are iteratively added to the region until the joint is completely encircled.

At each iteration, a cost function is evaluated for every node neighbouring the current region:

$$C^L = \sum_{k=1}^N \|x_k - x_{k-1}\| \times (1 + \|\alpha\|) \quad (1)$$

The cost function is a modification of Dijkstra’s shortest path algorithm which minimizes the length of a path between two nodes [8]. Parameter x refers to the position of a node. Parameter α prevents sharp inflections in the different branches and is described below.

Because the glenohumeral joint is a ball and socket joint, we can safely state that the joint separation loop will be located roughly in a single plane. To determine whether the direction of a growing tip of the path deviates from the direction of its predecessors, we evaluate the angle between the vector perpendicular to the front half of its preceding segments and the vector perpendicular to all of its preceding segments. Figure 4 illustrates how this angle α is determined.

Because we are interested in the joint separation loop, the paths should specifically connect nodes that have a negative minimum curvature. Therefore, when the minimum curvature of a particular path is positive for three consecutive steps, it is assumed the path has left the joint gap. The path is then discontinued. The number of three was chosen so that occasional bumps in the surface do not restrain the algorithm from finding a solution, while paths do not continue to expand outside of the joint gap.

The expectation is that a path evolves in two directions from a starting point and will eventually encircle the joint. In other words, the two tips of the path

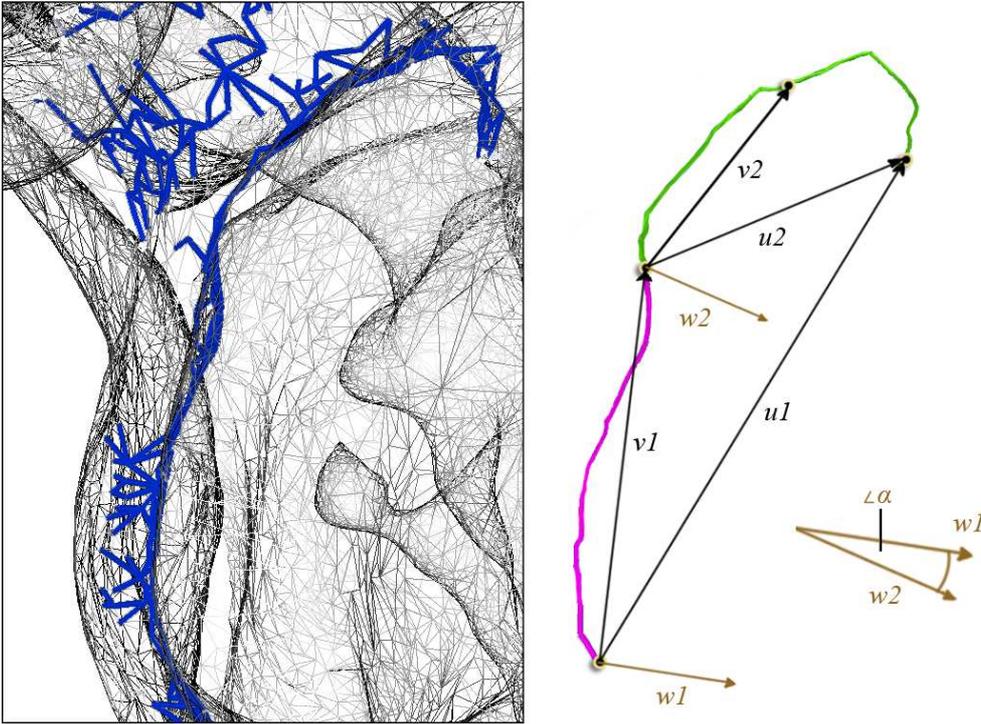


Figure 4: Parameter α of equation 1. The angle in radians is calculated between the vector perpendicular to the front half of the preceding segments (green) and the vector perpendicular to all of the preceding segments (green plus purple). The vectors w_1 and w_2 are directed perpendicular to v_1, u_1 and v_2, u_2 respectively. If the path takes sharp turns, the value of α will increase, thus increasing the total cost of the path. This prevents sharp inflections.

should meet each other. At each iteration, we check whether a newly added node neighbours any other part of the path. If so, this could indicate a successful encirclement. Two further conditions have to be satisfied:

1. The two tips have to run in opposite directions when they meet. This is the case if the dot product of their respective directions is negative.
2. The total path has to form a closed loop. This is the case if the accumulated path normal is close to 0. The accumulated path normal is defined as the sum of all surface normals crossed by the path, each scaled by the inverse of the sum of the edges before and after the vertex that it is defined on.

When the algorithm terminates, the result is a closed loop that encircles the glenohumeral joint. We refer to this loop as the joint separation loop.

2.3. *Slice-by-slice Separation Contour*

In this stage the joint separation loop is used to derive, for each slice of the unsmoothed isotropic volume (see Section 2.1 Preprocessing), a 2-D contour separating humerus and scapula. Together, these contours will form a surface that separates the humerus and scapula in 3-D. On each slice, the two intersection points of the joint separation loop with that slice serve as the end-points of the 2-D separation contour.

In the case of a healthy joint, the humerus and scapula would be separated by a region of low intensity voxels. In our case, the more dense outer layers of the humerus and scapula are pressed together, resulting in a transition region of more or less consistent, but not necessarily low, density. The 2-D separation contour attempts to connect its endpoints whilst remaining within this transition region and crossing pixels of similar values to the values of the endpoints.

To connect a pair of endpoints a path seeking algorithm is used. The cost function attempts to minimise the variation between consecutive pixels in the path. In our data, bone has a HU of at least 100, while soft tissue surrounding the bone has a range somewhere around 0 HU. To capture this variation, we determine the absolute difference between the HU of a pixel and the HU of its predecesing pixel, divided by 100 and clipped to [0.0, 1.0]. These last two steps are required to control the influence of the variation within the cost function, relative to the additional parameter that is introduced in Section 2.5.

The cost function is defined as:

$$C^H = \sum_{k=1}^N \min\left(1, \frac{||H_k - H_{k-1}||}{100}\right) \quad (2)$$

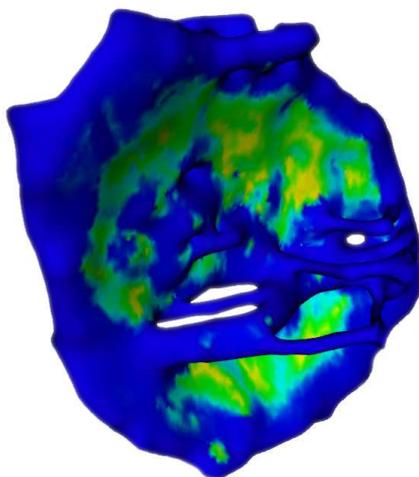


Figure 5: Surface model of the stacked 2-D separation contours. Highlighted areas indicate a high confidence value γ , further explained in Section 2.4.

where H is the HU of a pixel. From both points a path evolves that minimizes the cost function. When the two paths find each other, the path seeking algorithm terminates and reinitializes for the point pair of the next slice.

It is highly unlikely that the segmentation is correct after the initial pass, because the pathological state of the shoulder allows the paths to run through cortical bone, the thin outer layer of bone that normally has a high density and strength. However, even an erroneous segmentation provides a useful basis for the subsequent refinement stage.

2.4. Hough Feature Volume

In this stage the 2-D separation contours are used to derive a feature volume, based on the Hough-transform, that will be used in the subsequent stage to extract a surface that accurately separates the humerus and the scapula.

An encapsulating surface is constructed of the rasterised separation contours (see Figure 5). The surface is Laplacian smoothed with 20 iterations. This method yields suitable surface normals for further analysis. In spite of this, it would be interesting to investigate alternative volume-preserving smoothing methods in the future, such as those proposed by Taubin [23] and Vollmer et al.[26].

We apply a Hough transform to the smoothed surface [9]. For each point on the smoothed surface we follow its extended normal through an empty volume,

rasterizing it using Bresenham’s algorithm [5]. The result is a volume where each voxel contains a scalar equal to the number of times that an extended surface normal of the smoothed surface has intersected that particular voxel. Due to the sphericity of the humeral head and hence the smoothed surface, the area under the center of the head contains consistently higher values.

Again, for each point on the smoothed surface we follow its extended normal through the Hough volume and determine the position of the voxel that contains the highest number of intersections. This position is referred to as the point origin of the surface point that it corresponds with. The confidence value γ of a surface point is defined as the ratio between the Hough value at its point origin and the maximal value of the Hough volume. In other words, if a point origin is located at a globally high Hough value, it is likely that the surface point is located within the joint space and therefore we assign a high confidence. For each smoothed surface point, the point origin, the vector to the point origin and the confidence are stored.

The smoothed surface is voxelized using 3-D splatting: For each surface point, a Gaussian kernel (standard deviation 1.5, radius 5 voxels), weighted by the confidence γ for the point, is centered on the closest voxel position for that point. The point’s stored information is additively spread to surrounding voxels via the confidence weighted kernel. After all surface points have been traversed, the voxelisation is normalized with the per-voxel accumulated confidence. The per-voxel accumulated confidence is separately normalised by the maximum per-voxel accumulated confidence.

For every voxel an agreement value λ is determined, defined as the extent to which the voxel location is in agreement with the Hough surface determined by the neighbouring voxels (see Figure 6).

2.5. Iterative Refinement

In this step, the slice-by-slice separation contours are iteratively recalculated in accordance with the Hough feature volume. The resulting slice-by-slice separation contours form the final segmentation.

As shown in Figure 2, after their initial execution, stages 2.2 Joint Separation Loop and 2.3 Slice-by-slice Separation Contour are repeated until a sufficiently accurate set of separation contours are produced. With each iteration, new contours are calculated based on the Hough feature volume of the previous iteration, and based on these a new Hough feature volume is derived. The cost function C^T , shown in equation 3 and used in all repeated iterations is a modified version of the cost function C^H used in the initial iteration. Whereas C^H takes into account

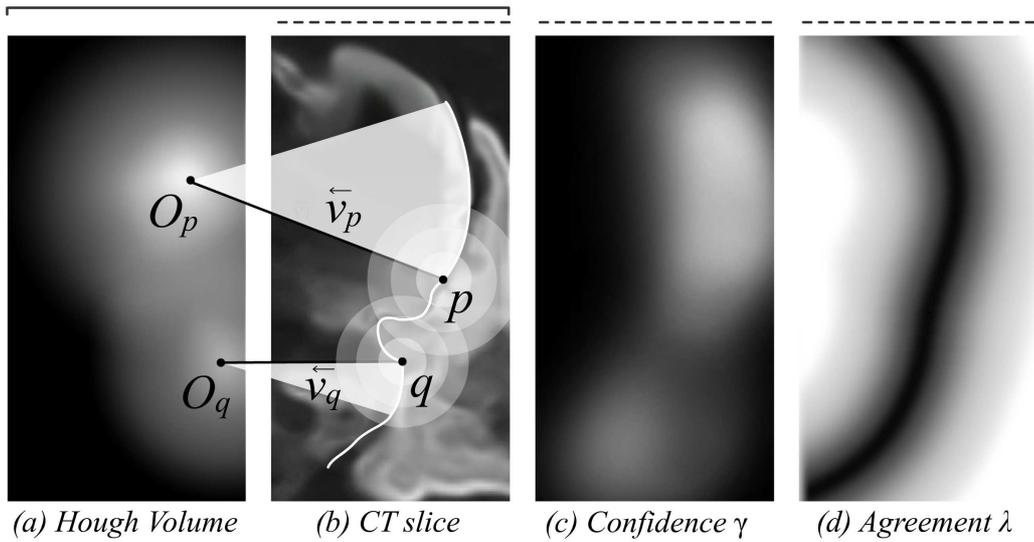


Figure 6: (a) The Hough volume is determined by applying a Hough transform to the surface formed by the slice-by-slice separation contours. (b) For each point on this surface a point origin O is determined, together with a vector \vec{v} pointing to this point. (c) Confidence γ is a fraction of the Hough value of the surface point's point origin, relative to the maximum value in the Hough volume. The values of O , \vec{v} and γ are 3D-splatted across the surface using Gaussian kernels. For all voxels we determine the distance to their 3D-splatted point origins. (d) The agreement λ is the extent to which this distance and the length of the 3D-splatted vector \vec{v} correspond with one another.

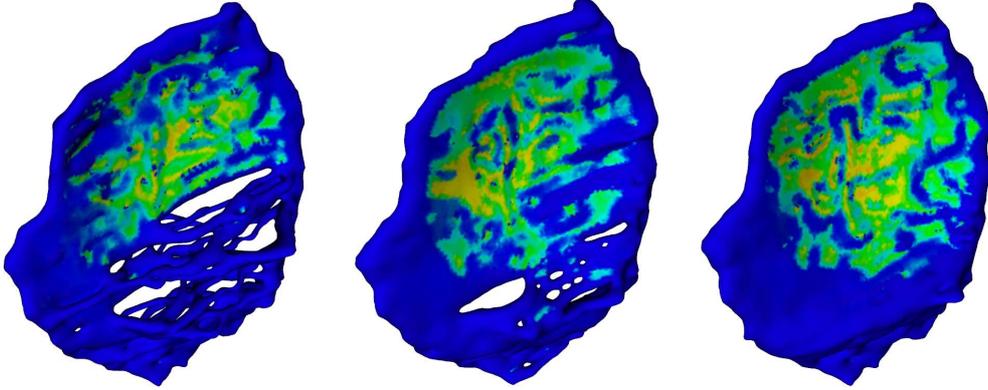


Figure 7: Three iterations of the segmentation algorithm. The color mapping equals that of Figure 5. With each iteration, the smoothly curved surfaces influence the cost-function of neighbouring areas by evaluating whether a separation contour is in agreement with nearby parts of the Hough surface.

only the variations in volume intensity along the contour, C^T takes into account the agreement value λ described in Section 2.4 as well.

$$C^T = \sum_{k=1}^N \gamma_k \times \lambda_k + (1 - \gamma_k) \times \min\left(1, \frac{\|H_k - H_{k-1}\|}{100}\right) \quad (3)$$

More specifically, if confidence γ is low, i.e. a point is not close to the Hough surface, its cost is primarily determined by its intensity variation. However, if confidence γ is high, i.e. the point is close to the Hough surface, the agreement λ of a point contributes more heavily to the cost: High agreement leads to low cost and vice versa. With this, we ensure that contours run through surface points that are located close to the Hough surface and are in agreement with their neighbours as to the location of that surface.

The algorithm is iterated until convergence. The criterion for this is that two consecutive steps produce the same segmentation result. Inspection showed that for most of the datasets four iterations of the algorithm are sufficient for convergence. After the last iteration the joint separation lines are used to separate the two bones and create two distinct volume masks.

2.6. Evaluation

We have evaluated our technique on ten shoulder CT-datasets. The datasets we used for testing were acquired from the hospital PACS and were performed

Dataset	X-Ray Tube Current	Exposure Time	Spatial Resolution
1	120	1000	$0.900 \times 0.900 \times 1.000$
2	160	500	$0.900 \times 0.900 \times 1.000$
3	160	600	$0.858 \times 0.858 \times 1.000$
4	339	500	$0.970 \times 0.970 \times 1.000$
5	500	500	$0.488 \times 0.488 \times 1.000$
6	339	500	$0.970 \times 0.970 \times 1.000$
7	376	500	$0.885 \times 0.885 \times 2.000$
8	70	500	$0.412 \times 0.412 \times 1.000$
9	410	500	$0.934 \times 0.934 \times 2.000$
10	158	500	$0.919 \times 0.919 \times 1.000$

Table 1: Scan parameters of the evaluation scans.

over a period as part of the standard treatment workflow. As such, CT parameters vary (see Table 1). All shoulders were diagnosed by an orthopaedic surgeon and rated as the highest level (level 5) of rheumatoid arthritis using the Larsen-score.

3. Results

To determine the accuracy of the algorithm, we compared the resulting segmented voxel masks with voxel masks that were obtained with manual segmentation. The voxel masks consisted of separate humeral (upper-arm bone) and scapular (shoulder blade) volumes. Manual segmentation was performed by an expert orthopaedic surgeon. We extracted a volume of interest which contained the glenohumeral joint plus an additional 10% of its size in all directions.

For quantitative evaluation we calculated the Dice coefficient between all manual segmentations and the segmentations derived by our technique. The Dice coefficient P_{vo} , expressed as a percentage for convenience, is defined as follows:

$$P_{vo} = \frac{2 \times |S \cap R|}{|S| + |R|} \times 100\% \quad (4)$$

where S and R refer to the two segmented volumes that are being compared. A voxel-perfect segmentation results in a Dice coefficient of 100%.

Table 2 shows the Dice coefficient for all datasets for both the scapula and the humerus. For the humerus, the median Dice coefficient was 98.9% with an

Humerus		
Dataset	Dice coefficient [%]	Hausdorff distance [mm]
1	98.88	1.84
2	98.59	2.45
3	98.99	3.17
4	99.13	2.45
5	94.28	5.59
6	94.8	3.04
7	93.93	5.55
8	99.58	3.07
9	99.53	1.69
10	99.04	3.67

Scapula		
Dataset	Dice coefficient [%]	Hausdorff distance [mm]
1	98.28	2.11
2	98.6	4.52
3	99.18	9.28
4	98.34	1.52
5	98.47	4.75
6	97.05	2.97
7	97.84	4.83
8	99.53	9.19
9	99.53	1.26
10	99.25	3.30

Table 2: Validation results of 10 shoulder CT datasets. The first column shows the id of the CT dataset. The second column shows the Dice coefficient of the voxel mask as created by manual segmentations and as created by our technique. Column three shows the Hausdorff distances.

inter-quartile range (IQR) of 95.8% to 99.4%. For the scapula, the median Dice coefficient was 98.5% with an IQR of 98.3% to 99.2%.

The extent of the volumes was limited to the volume of interest as used for the segmentation steps, i.e. the joint gap. This entails that a small segmentation error will have a significant effect on the Dice coefficient. However, because we also wanted to have a qualitative measure independent of the size of the volume masks, we calculated the Hausdorff distances using Mesh 1.13 [1]. The Hausdorff distance is the maximum distance of a volume to the nearest point in the other volume and thus reflects the largest segmentation error. As a frame of reference, please note that the average diameter of proximal humeri is approximately 46 mm [3]. For the humerus, the median Hausdorff distance was 3.06 mm with an IQR of 2.30 mm to 4.14 mm. For the scapula, the median Hausdorff distance was 3.92 mm with an IQR of 1.96 mm to 5.92 mm. The Hausdorff distances were added to Table 2.

The results of an evaluation dataset together with its manual segmentation can be seen in Figure 8. The Dice coefficients for the initial segmentation and three iterations are 56.32%, 84.12%, 97.30% and 98.88% for the humerus model and 69.02%, 87.99%, 96.78% and 98.28% for the scapula model. These are typical increments of accuracy that we see for other datasets.

In all cases, the complete segmentation process completed in under two minutes on a 2 GHz Core T2500 laptop processor. The time needed to find the joint separation loop ranged from about five to about ten seconds.

4. Discussion

4.1. Limitations

For datasets 5 and 7 the humerus volume mask differed considerably from the ground truth segmentation. Also, the Hausdorff distances of the scapulae of datasets 3 and 8 were relatively large. Upon closer inspection we noticed that osteophytes, i.e. bone deformations, at the edge of the glenohumeral joint had been included in the automatically segmented volumes, while they had been excluded from the manual segmentations. Because the density of these osteophytes varies heavily, subtle segmentation differences may influence whether an osteophyte is included in the segmented volume. A possible improvement would be to highlight the osteophytes and allow the user to explicitly make these segmentation decisions.

Figure 8 and the corresponding Dice coefficients point out that the segmentation quality improves considerably during the first iterations. A side effect of our

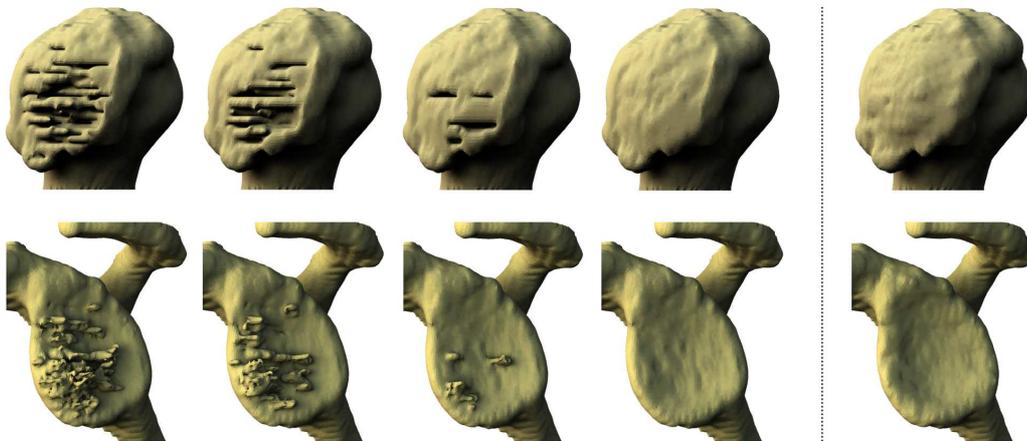


Figure 8: The scapula and humerus model of an evaluation dataset after the initial slice-by-slice segmentation and three subsequent iterations. The most right image shows the manual segmentation that was used to evaluate the quality of our segmentation.

approach is that when slice-by-slice separation contours of adjacent slices follow a smooth curve, the Hough feature volume will pick up this smooth area and force the slice-by-slice separation contours in subsequent iterations in the same erroneous direction, never improving the segmentation quality for these specific parts. In general, the erroneous parts of the slice-by-slice separation contours do not follow smooth curves, because they run through the pathological bone area rather than through the joint space. Consequently, this distortion has limited effect on the segmentation accuracy of our approach, as shown by the evaluation results.

The criteria used for determining and closing the joint separation loop perform very well for the evaluation datasets. It is conceivable that one of the criteria is not met, although we have not experienced this for any of the 10 evaluation datasets. In this case, no joint separation loop will be returned and the surgeon has to reselect the joint to retry. As demonstrated by Chambers et al. [6] the complexity of finding the shortest separating loop is a NP-hard problem. The criteria we use to find the joint separation loop (i.e. evaluation of the minimum curvature of the surface and prevention of sharp inflections) are useful exploits, reducing this problem to a routine that, at least for the 10 evaluation datasets we used, returns a joint separation loop within ten seconds. One can think of ways to improve this algorithm, for example by using more selection points, adding to the robustness of the algorithm.

Another limitation of this study that the quality of the evaluation data differed

due to the varying acquisition parameters and subject variability. For example, the contralateral shoulder of subject 6 contained a prosthesis, producing scattering effects in the CT-scan. Although the visible scatter in the segmented shoulder was limited, this additional noise may affect the segmentation routine. A future phantom study would be interesting to systematically assess the sensitivity of our algorithm to noise. However, the large variation in our data suggests that the technique is relatively robust and thus applicable to CT scans with different noise levels.

4.2. Conclusion

In this paper we have presented a novel segmentation technique that combines surface and volume processing to provide fast and accurate segmentation of arthritic glenohumeral joints from CT data. From the evaluation we conclude that our technique is sufficiently accurate for the segmentation of heavily deteriorated glenohumeral joints.

The segmentation results of our data collection of ten shoulders were compared to the manual segmentation as performed by an expert surgeon. The Dice coefficients and Hausdorff distances indicate that our technique yields highly accurate results compared to manual segmentation. Our technique was sufficiently robust to extract accurate segmentations from the datasets in spite of their pathological nature, as indicated by the Larsen-score level 5, the varying bone density and small joint space, and in spite of varying acquisition parameters.

To our knowledge, no other segmentation techniques exist that have been shown to cope with arthritic shoulder joints. As discussed in the introduction, Botha [2] has shown that level-sets and watershed segmentation have limited success, in the course of which setting the parameters is a time-consuming process not suitable for clinical practice. Although Zoroofi et al. [29] successfully segmented arthritic hip joints, they were less successful for arthritic hips with Larsen-score 4 and 5. The lack of fast segmentation techniques capable of segmenting arthritic shoulder joints was our motivation for the research described in this work. The technique is currently applied to a pre-operative planning application used in our clinic and will serve as a basis for surface-based orthopaedic applications that involve arthritic shoulder joints. The high accuracy together with the short time required for the segmentation process make this technique a good approach for the segmentation of glenohumeral joints in a clinical environment.

In future work we will test algorithm performance on other pathological joints, such as the hip and knee joint. Because these joints normally have a highly curved

surface like the glenohumeral joint, we expect that the algorithm may also work on these joints.

Acknowledgements

The authors wish to thank Lingxiao Zhao for his assistance in this work.

References

- [1] Aspert, N., Santa-Cruz, D., Ebrahimi, T., 2002. Mesh: Measuring errors between surfaces using the hausdorff distance. In: Proceedings of the IEEE International Conference on Multimedia and Expo. Vol. I. pp. 705 – 708.
- [2] Botha, C. P., 2005. Techniques and software architectures for medical visualisation and image processing. Ph.D. thesis, Delft University of Technology. URL: <http://visualisation.tudelft.nl/People/CharlBotha/PhDThesis>
- [3] Boileau, P., Walch, G., Sep 1997. The three-dimensional geometry of the proximal humerus. implications for surgical technique and prosthetic design. *J Bone Joint Surg Br* 79 (5), 857–865.
- [4] Branzan-Albu, A., Laurendeau, D., Hbert, L., Moffet, H., Dufour, M., Moisan, C., June 17-18 2004. Image-guided analysis of shoulder pathologies: Modeling the 3d deformation of the subacromial space during arm flexion and abduction. In: Stephane Cotin, D. M. (Ed.), International Symposium on Medical Simulation (ISMS 2004). Vol. 1 of Lecture Notes in Computer Science. Springer Verlag, Cambridge, MA, USA, pp. 193–202.
- [5] Bresenham, J., January 1965. Algorithm for computer control of a digital plotter. *IBM Systems Journal* 4 (4), 25–30.
- [6] Chambers, E. W., ric Colin de Verdire, Erickson, J., Lazarus, F., Whittlesey, K., 2008. Splitting (complicated) surfaces is hard. *Computational Geometry* 41 (1-2), 94 – 110, special Issue on the 22nd European Workshop on Computational Geometry (EuroCG), 22nd European Workshop on Computational Geometry.
- [7] Di Gioia, III, M. A. M., Blendea, S., Jaramaz, B., 2007. Surgical navigation for total hip arthroplasty. In: J.J. Callaghan, A.G. Rosenberg, H. R. (Ed.), *The Adult Hip*, 2nd Edition. Vol. 2. Lippincott Williams & Wilkins, pp. 1053 – 1059.

- [8] Dijkstra, E., 1959. A note on two problems in connexion with graphs. *Numerische Mathematik* 1 (1), 269–271.
- [9] Duda, R., Hart, P., 1972. Use of the Hough transform to detect lines and curves in pictures *Comm. In: ACM*, 1972, 15, 11–15.
- [10] Fleute, M., Lavallee, S., 1998. Building a complete surface model from sparse data using statistical shape models: application to computer assisted knee surgery system. In: *MICCAI*. pp. 879–887.
- [11] Kang, Y., Engelke, K., Kalender, W. A., 2003. A new accurate and precise 3-d segmentation method for skeletal structures in volumetric ct data. *IEEE Transactions on Medical Imaging* 22 (5), 586–598.
- [12] Krekel, P., Botha, C., Valstar, E., de Bruin, P., Rozing, P., Post, F., March 2006. Interactive simulation and comparative visualisation of the bone-determined range of motion of the human shoulder. In: Schulze, T., Horton, G., Preim, B., Schlechtweg, S. (Eds.), *Proc of Simulation and Visualization*. SCS Publishing House Erlangen, pp. 275–288, best Paper Award.
- [13] Larsen, A., 1995. How to apply larsen score in evaluating radiographs of rheumatoid arthritis in long-term studies. *J Rheumatol* 22, 19745.
- [14] Lorensen, W. E., Cline, H. E., 1987. Marching cubes: A high resolution 3d surface construction algorithm. *Computer Graphics* 21 (4).
- [15] Mangan, A. P., Whitaker, R. T., 1999. Partitioning 3d surface meshes using watershed segmentation. *IEEE Transactions on Visualization and Computer Graphics* 5 (4), 308–321.
- [16] Page, D. L., Sun, Y., Koschan, A. F., Paik, Paik, J., Abidi, M. A., 2002. Normal vector voting: crease detection and curvature estimation on large, noisy meshes. *Graphical Models* 64.
- [17] Page, D. L., Koschan, A. F., Abidi, M. A., 2003. Perception-based 3d triangle mesh segmentation using fast marching watersheds. *cvpr 02*, 27.
- [18] Pitiot, A., Delingette, H., Thompson, P. M., Ayache, N., 2004. Expert knowledge-guided segmentation system for brain mri. *Neuroimage* 23 Suppl 1, S85–S96.

- [19] Prescher, A., Aug 2000. Anatomical basics, variations, and degenerative changes of the shoulder joint and shoulder girdle. *Eur J Radiol* 35 (2), 88–102.
- [20] Rajamani, K. T., Styner, M. A., Talib, H., Zheng, G., Nolte, L. P., Ballester, M. A. G., Apr 2007. Statistical deformable bone models for robust 3d surface extrapolation from sparse data. *Med Image Anal* 11 (2), 99–109.
- [21] Sato, Y., Sasama, T., Sugano, N., Nakahodo, K., Nishii, T., Ozono, K., Yonenobu, K., Ochi, T., Tamura, S., 2000. Intraoperative simulation and planning using a combined acetabular and femoral (caf) navigation system for total hip replacement. pp. 1114–1125.
- [22] Tao, X., Prince, J. L., Davatzikos, C., May 2002. Using a statistical shape model to extract sulcal curves on the outer cortex of the human brain. *IEEE Trans Med Imaging* 21 (5), 513–524.
- [23] Taubin, G., 1995. A signal processing approach to fair surface design. In: *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, New York, NY, USA, pp. 351–358.
- [24] Tremblay, M.-E., Branzan-Albu, A., Laurendeau, D., Hbert, L., May 17-19 2004. Integrating region and edge information for the automatic segmentation for interventional magnetic resonance images of the shoulder complex. In: *1st Canadian Conference on Computer and Robot Vision (CRV2004)*. Vol. 1. IEEE Computer Society, London, Ontario, Canada, pp. 279–286.
- [25] van der Glas, M., Vos, F. M., Botha, C. P., Vossepoel, A. M., 2002. Determination of Position and Radius of Ball Joints. In: Sonka, M. (Ed.), *Proceedings of the SPIE International Symposium on Medical Imaging*. Vol. 4684 - Image Processing.
- [26] Vollmer, J., Mencil, R., Mueller, H., 1999. Improved laplacian smoothing of noisy surface meshes. In: *Computer Graphics Forum*. Vol. 18. pp. 131–138.
- [27] von Schroeder, H. P., Kuiper, S. D., Botte, M. J., Feb 2001. Osseous anatomy of the scapula. *Clin Orthop Relat Res* (383), 131–139.

- [28] Wu, C.-H., Sun, Y.-N., Dec 2006. Segmentation of kidney from ultrasound b-mode images with texture-based classification. *Comput Methods Programs Biomed* 84 (2-3), 114–123.
- [29] Zoroofi, R. A., Sato, Y., Sasama, T., Nishii, T., Sugano, N., Yonenobu, K., Yoshikawa, H., Ochi, T., Tamura, S., 2003. Automated segmentation of acetabulum and femoral head from 3-d ct images. *IEEE Transactions on Information Technology in Biomedicine* 7 (4), 329–343.